

F. R. Bailey  
NASA Ames Research Center  
Moffett Field, California

### INTRODUCTION

The Numerical Aerodynamic Simulation (NAS) Program had its inception in 1975 at the Ames Research Center (ARC) when a small group of researchers associated with the computational fluid dynamics program set out to obtain significantly greater computer power and the memory capacity needed to solve three-dimensional fluid flow models (Peterson et al., 1984). The state of the art in computational aerodynamics was at the point where problems involving complex geometries could be treated only with very simple physical models and only those involving simple two-dimensional geometry could be treated with more complex physical models. It was clear that to treat problems with both complex three-dimensional geometries and complex physics required more computer power and memory capacity than was available. This is illustrated in figure 1, where the estimated computer speed and memory requirements

derived from assuming a solution in 15 min of central-processor time are compared with the capabilities of several supercomputer generations (Peterson et al., 1985). As indicated, current supercomputers can adequately address inviscid flows, but the computers needed to adequately address more complex Reynolds-averaged approximations to turbulent flows about full aircraft configurations will not be available until the end of this decade. Large-eddy simulations for complete aircraft must wait for future, more powerful computers.

Sparked by this need for increased computer power, the ARC team spent several years performing requirements refinement, technical studies, and advocacy that resulted in the establishment of the NAS Program in 1983. The ongoing objectives of the program are: (1) to provide a national computational capability, available to NASA, Department of Defense, and other Government agencies, industry, and universities, as a necessary element in




Fig. 1 Computer speed and memory requirements for aerodynamic calculations compared with the capability of various supercomputers.

ensuring continuing leadership in computational fluid dynamics and related disciplines; (2) to act as a pathfinder in advanced, large-scale computer system capability through systematic incorporation of state-of-the-art improvements in computer hardware and software technologies; and (3) to provide a strong research tool for NASA's Office of Aeronautics and Space Technology.

An early task in implementing the NAS Program was to ensure its role as a pathfinder in providing advanced supercomputing capabilities to the computational fluid dynamics community. To accomplish this a long-term strategy of installing, at the earliest opportunity, a system representing each new generation of increasingly more powerful high-speed processors was initiated. Each generation may be a prototype or early production model. The strategy, illustrated in figure 2, has already started with the installation of Ser. No. 2 of the Cray-2 as the first generation, high-speed processor. Introduction of the Cray-2 required an 8-mo period of test and integration before production operations began. The second high-speed processor will be added in 1987 followed by several months of test and integration leading to an operational capability in 1988. The third generation will replace the first in the 1989/90 time frame. In subsequent years, new advanced high-speed processors will replace the older installed systems, thereby maintaining a steady state in which there are two high-speed processors, one of which is the most advanced available.

Coupled with the high-speed processor installation strategy, is the ongoing design and development of the NAS Processing System Network (NPSN). The NPSN is a network of high-speed processors and support computers configured to provide a state-of-the-art, scientific, supercomputing environment to the computational fluid

dynamics and computational physics community. The NPSN will be implemented in phases, with the completion of each phase resulting in an NPSN configuration keyed to the introduction of a new-generation, high-speed processor. The NAS Program defined the first two configurations in 1983. These are the Initial Operating Configuration (IOC) and Extended Operating Configuration (EOC) defined below.

#### Initial Operating Configuration:

- High-speed processor-1 (HSP-1)
  - 250 million floating-point operations/sec (MFLOPS)
  - computing rate for optimized, large-scale, computational aerodynamic applications
  - 256 million 64-bit words of central memory
- Integrated and expandable NPSN configuration
- Common-user interface for all subsystems
- Medium-band (56 kilobits/sec) communications for remote sites
- Wide-band (1.544 megabits/sec) communications to NASA centers

#### Extended Operating Configuration:

- Additional High-Speed Processor-2 (HSP-2)
  - 1000 MFLOPS computing rate for optimized, large-scale, computational aerodynamic applications
- Upgraded subsystems to accommodate HSP-2 and graphics subsystems
- 6.2 Megabits/sec communications to NASA centers

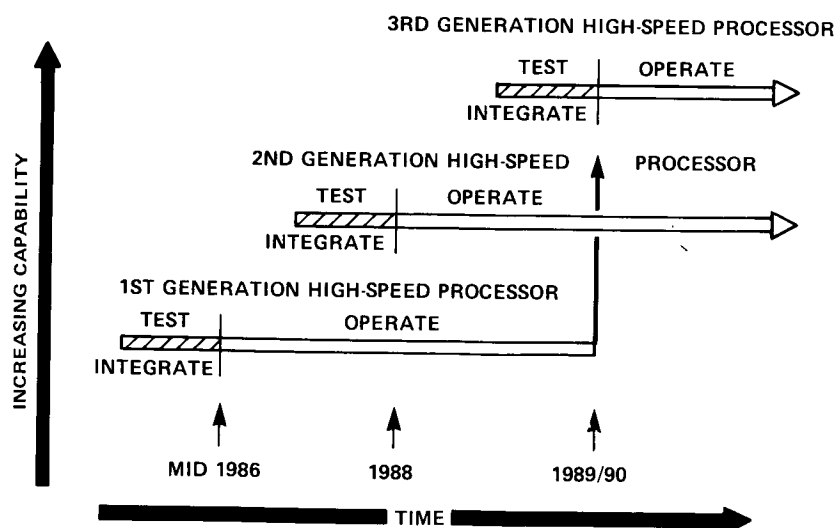


Fig. 2 NAS Program strategy for introducing state-of-the-art high-speed processors.

The NAS Facility, completed in January 1987, provides the physical location for the NAS Program. It houses the NPSN, system support and development staff, and computational fluid dynamics research staff. This new 90,000-ft<sup>2</sup> building was designed with approximately 15,000 ft<sup>2</sup> of raised computer floor which is sufficient to accommodate future expanded NPSN configurations. Approximately 15,000 ft<sup>2</sup> of additional raised floor is available for technical support space, laboratories, and user colocated workstation equipment. Office space was designed to allow the computer systems staff and computational fluid dynamics researchers to share the building. These accommodations provide a strong bonding between the computer service providers and their researcher clients.

The NAS Systems Division, formerly the NAS Projects Office, at ARC has completed the initial tasks and is responsible for accomplishing the NAS Program long-term objectives on a continuing basis. The Division includes a skilled staff of computer scientists and engineers who are responsible for planning, designing, and implementing the advanced technology necessary to keep NAS attuned to its pathfinding role. The Division also provides day-to-day computer operations; ongoing computer system support functions such as hardware and software maintenance; and logistics; as well as computer services personnel to provide the training, consulting and trouble-shooting necessary to offer premier supercomputing service to a national user base.

## NAS IOC DESCRIPTION

The IOC of the NPSN consists of the following functional subsystems: high-speed processor, support processing, workstation, mass-storage, and long-haul communications. The configuration, shown in figure 3, is a local-area-network-oriented architecture consisting of hardware from a number of vendors. To reduce hardware and vendor-specific dependencies, the NPSN software presents a common user interface across the system that provides users with the same environment on all user-accessible machines, i.e., provides common utilities and commands on user-visible subsystems. This uniform environment enables users to move easily between processors, allows for easy file access and command initiation across machine boundaries, and enhances user-code portability within the NPSN.

The operating system on all user-visible computers is based on AT&T's UNIX™ System V.2 operating system with network software modeled after the Berkeley 4.2/4.3 bsd UNIX operating system. Jobs on all user-visible systems can be run in batch or interactive modes. Communications protocols come from the DoD Internet family of protocols, commonly referred to as TCP/IP.

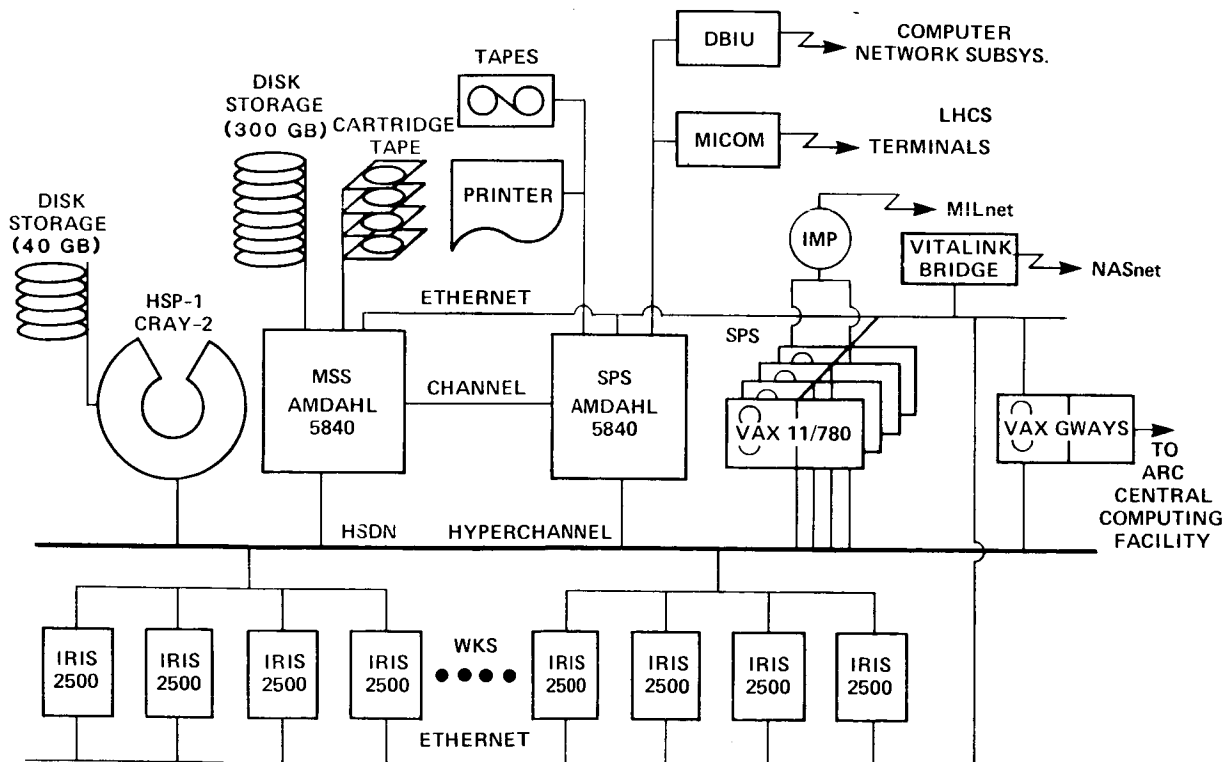


Fig. 3 IOC of the NPSN.

The first HSP, a Cray-2, is a four-processor system with 268,435,456 words of dynamic memory--or as more commonly referenced, with 256 megawords (where each 1k equals 1024 64 bits/word) of dynamic memory. The four processors can operate independently on separate jobs or in a multiprocessing mode on a single job. The NAS Cray-2 has achieved a peak computing speed of 1720 MFLOPS for assembly-language matrix multiplies and a sustained aggregate-speed in excess of 250 MFLOPS for a job mix of computational physics FORTRAN codes.

Common memory is accessed automatically by the hardware, and can be accessed randomly from any of the four processors as well as from any of the common data channels. Any job can utilize all or part of the common memory. This memory is divided into four quadrants, each of which has 128 interleaved banks. In addition to the common memory, each of the processors has a very-high-speed local memory for temporary storage of vector and scalar data. Externally, there are 34 DD49 disks attached to the Cray-2, giving a combined storage capacity of approximately 40 gigabytes. Cray-2 software includes an operating system (UNICOS) based on AT&T UNIX System V.2; an automatic vectorizing FORTRAN compiler; a C-language compiler; a large set of batch and interactive utilities; a large set of libraries, including a multitasking library; TCP/IP networking; and Berkeley "r" commands.

#### Support Processing Subsystem

The support processing subsystem (SPS) consists of an Amdahl 5860 mainframe computer and four VAX 11/780 minicomputers. This subsystem supports general-purpose interactive processing for local and remote users. In addition, it provides the interface to local and remote terminals and to tapes and high-speed printers.

#### Workstation Subsystem

The Workstation Subsystem (WKS) consists of microprocessor-based workstations which produce graphical displays of Cray-2-generated datasets and manipulate text and data as well as perform small-scale computations. The IRIS™ (Integrated Raster Imaging System), selected as the initial NAS workstation, is manufactured by Silicon Graphics Inc. and interfaces to both Ethernet™ and Network Systems Corporation's HYPERchannel™ networks. A key element of the IRIS is a set of special-purpose graphics microprocessors, called a "geometry engine," which transforms and displays three-dimensional data at rates exceeding 50,000 floating point coordinates/sec. Other workstations with TCP/IP interface capabilities are also compatible with the NPSN System.

The Mass Storage Subsystem (MSS) consists of an Amdahl 5860 mainframe computer, Amdahl 6380 disks, and IBM 3480 cartridge tape drives; this subsystem serves the global online and archival storage needs of the NPSN. The MSS checks and coordinates requests for files to be stored or retrieved within the NPSN and maintains a directory of all contained files. The MSS also acts as a file server for the other NPSN subsystems, controls its own internal devices, and performs file duplication, media migration, storage allocation, accounting and file management. The aggregate Amdahl 6380 disk-storage capacity is approximately 300 gigabytes at present.

#### High-Speed Data Network

The High-Speed Data Network (HSDN) is the local communications network which allows the exchange of data and control messages within the NPSN. Major design emphasis was placed on its ability to support large file transfers between the subsystems. The local network includes HYPERchannel, Ethernet, and the driver-level network software to support the intra-NAS data communications.

#### Long-Haul Communication Subsystem

NPSN access by the nationwide remote user community is provided by the Long-Haul Communication Subsystem (LHCS). LHCS consists of a NAS-unique remote-connection network, NASnet, and hardware/software interfaces to existing Government-sponsored networks.

NASnet, shown in figure 4, is a starlike configuration of data communication links connecting remote user sites to the NPSN. The physical data links between NASA Centers are provided by NASA's Program Support Communication Network. These links are extended to non-NASA sites via AT&T tail circuits. Except for dedicated 224 Kbit/sec terrestrial links to Lewis and Langley Research Centers, all data links are switched terrestrial circuits with a capacity of 56 Kbit/sec. At each termination point there is a Vitalink TransLAN™ communication bridge that connects the data circuit to an Ethernet network. The bridges permit NPSN Ethernet packets with remote addresses to be sent to destination hosts on the remote-site Ethernets. Conversely, Ethernet packets on a remote site's Ethernet can be routed to NPSN Ethernet hosts. In this way, remote users see their own local host, i.e., workstation, as just another node on the NPSN Ethernet. Since NASnet is implemented by switch circuit technology, the remote user literally dials up the connection only when required. From the software viewpoint, the selection of the DoD Internet Protocol (TCP/IP) suite for the NPSN

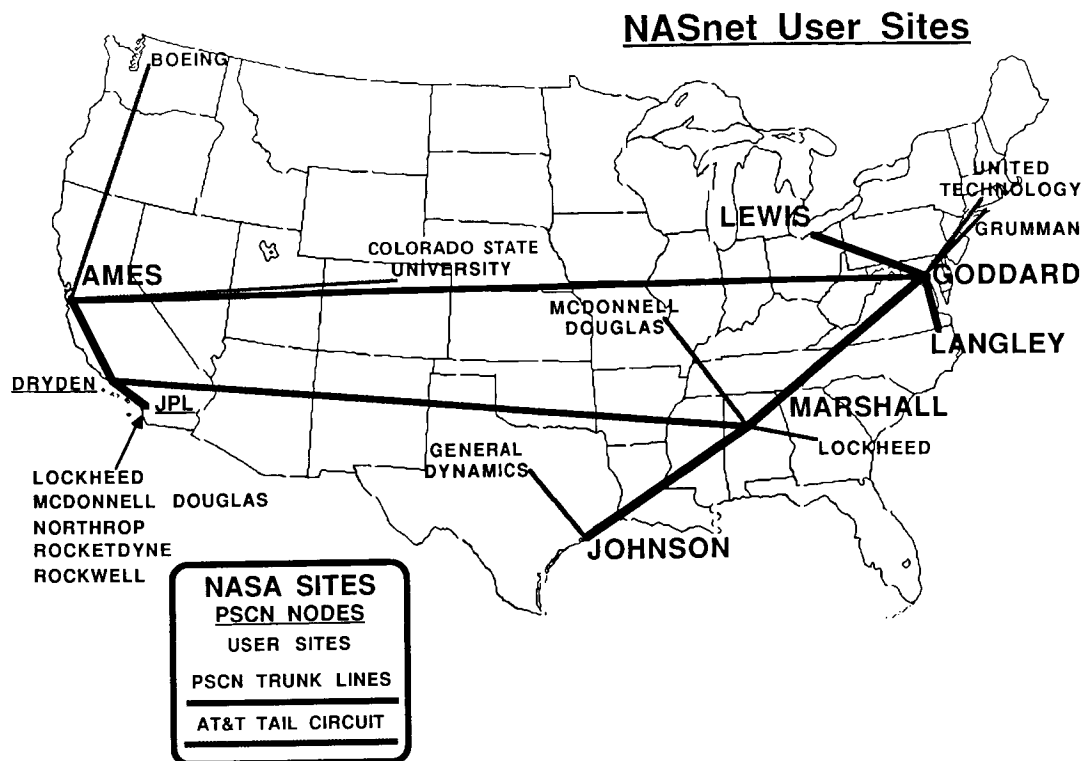


Fig. 4 NASnet configuration.

allows the most straightforward implementation of NASnet since it not only supports a rich set of applications (mail, file transfer, remote login, etc.), but implementation also exists for a large number of various computers ranging from IBM PCs to Cray supercomputers. In summary, NASnet, whose first prototype implementation was in July 1985, provides the remote user with all the functional capabilities of a local user of the NPSN and, in particular, gives the ability to use all the interactive capabilities of powerful graphics workstations. For example, experience has shown NASnet to be very effective in supporting NAS workstation clusters at Lewis and Langley Research Centers. Users at these Centers not only access the NPSN in the same manner as ARC users, but also utilize their Ethernet connected workstations to access computer resources at their own Centers.

In addition to NASnet, PSCN circuits support two other long-haul communication capabilities. The first is the Computer Network Subsystem (CNS) that provides batch-oriented bulk file transfer service among NASA Office of Aeronautics and Space

Technology (OAST) Research Centers (Ames-Moffett, Ames-Dryden, Lewis and Langley) using T1 (1.544 mbit/sec) satellite links.

The second PSCN-provided capability is the NASA Packet Switch System (NPSS). NPSS provides 9600 bit/sec terminal access to the NPSN from NASA Centers. Since NPSS is also connected to GTE's Telenet™ packet switched network, remote users outside NASA Centers can gain access to the NPSN at rates of 1200-2400 bits/sec.

Access by DoD research installations and many universities is provided by MILnet and ARPAnet. Additional university access will be provided in the future by a newly formed network sponsored by the National Service Foundation (NSF) (Jennings et al., 1986). This network (popularly called NSFnet) is designed to provide remote access to NSF-sponsored supercomputer centers by connecting several regional networks to a national backbone circuit. One of these regional networks is the NSF-funded Bay Area Regional Research Network (BARRnet) which includes ARC, University of

California campuses (Berkeley, Davis, San Francisco, and Santa Cruz), and Stanford University. When operational in 1987, BARRnet will provide T1 communications links between participating member sites and will connect to ARPAnet at Stanford and to the NSF funded San Diego Supercomputer Center (SDSC) through a remote user access computer (RUAC) at UC Berkeley.

#### NAS COMPUTATIONAL SERVICES

The NAS Program provides all users with consulting, training, and documentation services. A telephone hotline is manned around the clock, excluding Government holidays. Consulting advice is also available by computer mail and in person (for local and visiting scientists). Consulting expertise includes NPSN operating systems, compilers, communications, code conversion, code optimization, file handling, tape usage, workstation graphics, and other areas of user interest.

Training services include 3-day user training classes which are held monthly at ARC. These classes cover introductory UNIX, code conversion aids, code optimization, file archival methods, tape usage, and user responsibilities. User training classes can be presented at remote sites by special arrangement. NPSN usage information is also available online in the form of manual pages as well as tutorials and sample sessions.

In the area of documentation, the NAS User Guide is sent to each account holder. The User Guide presents basic, high-level information. Other user documents are referenced for more in-depth information. The User Guide contains a complete list of Cray-2 manuals, editor manuals, reference cards, etc., which are available by request for all users.

In addition, NAS System Bulletins are sent to each user when major systems changes occur, and a monthly NAS newsletter is sent to users and other interested people. The newsletters contain NAS usage information along with items about the user projects and general NAS news and developments.

#### NAS USAGE

The NAS Program is intended to support pioneering work in basic and applied research. Approximately 90% of the scientific research pertains to the areas of fluid dynamics, aerodynamics, structures, material science, aerothermal dynamics, and other aeronautics applications. The remaining 10% includes work in other disciplines such as chemistry, atmospheric modeling, astrophysics, and other areas of interest to NASA. About 55% of the available computer resources supports NASA programs and is used by both in-house staff and organizations involved in NASA

grants, contracts, and joint programs. The approximate allocation of the remaining amount is 20% to DoD, 15% to industry, 5% to other Government agencies, and 5% to universities. Only the costs of industry proprietary work are reimbursed to the Government in accordance with a formula that includes operating and capital improvement costs.

The NAS Program has implemented project selection and resource allocation procedures which invite applications for NAS usage from the scientific user community and provide for peer review of projects and equitable resource allocation. The NASA Office of Aeronautics and Space Technology receives approximately 45% of the resources, to be shared among Ames, Langley, and Lewis Research Centers. Each Center suballocates its share to individual in-house and supported research projects. The university allocation is administered by the National Science Foundation. Overall resource availability is determined by the NAS Systems Division.

To be selected for NAS usage, potential users must propose technically sound and relevant projects which require the unique capabilities offered by the NPSN. The NAS selection criteria are: technical quality, national need, timeliness, and suitability to NAS resources. Selection criteria and allocation guidelines are periodically reviewed to ensure maximum benefits. The guidelines for usage described here apply only to NAS computers in an operational status. When new computers are being integrated into the NPSN, they are devoted primarily to testing and software development.

During the interim operational period, from July 1986 through February 1987, Cray-2 computer time was allocated to 123 projects. The selection process is now under way for the first full operational period, which is for 1 year, March 1987 through February 1988. For the interim period, Cray-2 time was distributed over a wide range of disciplines as shown in figure 5. Note that a high percentage of the Cray-2 was allocated to projects related to the National AeroSpace Plane (NASP). This reflects the importance of supercomputing to this new initiative. NASP technology applications are expected to increase in future years.

#### EARLY NAS OPERATIONAL EXPERIENCE

The NAS Program initiated pilot operations for a select group of users in December 1985, shortly after the Cray-2 had completed acceptance testing. In July 1986, interim operations of a completely integrated NPSN began. User access during these periods increased rapidly and appreciable operational experience was gained by NAS users.




Fig. 5 Cray-2 computer time allocation distribution based on 123 interim operational period allocations.

Even within the context of pilot operations, which included a high level of system development activities, initial NAS users provided overwhelmingly positive feedback on their increased productivity while using the NPSN. In particular, most users indicated a significant reduction in code development/debugging time, start-up time for changes in configuration geometry or grid resolution, and job turnaround times, as well as improvements in the graphical analysis of computed results.

The reduction in code development and debugging time can be attributed mainly to the availability of interactive supercomputing and the rich variety of tools in the UNIX environment. Prior to the development of the UNICOS operating system, vendor-supplied operating systems for supercomputers had remained, almost exclusively, as batch mode processing. While batch processing may be desirable for long production runs, debugging and code development are not best handled in this mode. UNIX provides both the interactive capability necessary for efficient code development and debugging, and the batch capability for production runs via the Network Queuing System (NQS) software. For the system developer, many of the same advantages are realized. In addition, it has been reported that the introduction of new systems, as well as software enhancements and "bugfixes," are more easily and quickly completed with the uniform UNIX environment and consistent protocol suite across the NAS machines.

The NAS Cray-2 is the first full-memory, four-processor Cray-2. The 256-megaword central memory capacity represents a two-orders-of-magnitude increase over previously available capacity. Prior to the Cray-2, increased resolu-

tion or geometric complexity required the use of a Solid State Disk (SSD). The input/output (I/O) delays associated with paging data between the central memory and disk were significant and, not infrequently, reached 10 or 20 times the actual CPU time for the job. Therefore, specialized, machine-specific code modifications were required for efficient utilization of SSD memory.

For example, a typical three-dimensional Reynolds-averaged Navier-Stokes simulation is constrained to resolve less than 69 cubed (approximately 0.3 megaword) spatial points if the job is to remain in the central memory of an 8-megaword machine. Full use of 256 megawords would allow for 217 cubed, or 10.24 million points, to be resolved for an in-core run of this example code. Thus the increased memory of the Cray-2 allows for both increased job complexity without significant code modification and virtual elimination of I/O wait time. The net effect is decreased startup time for increasingly complex problems and improved job turnaround.

As the speed and memory of supercomputers continue to increase, larger and more physically complex problems can be solved more routinely. To analyze these larger and more complex datasets, graphical displays are not only desirable but necessary. Specialized graphical processors such as the "geometry engine" of the IRIS workstation help to provide pseudo-real-time visualization of computed results, i.e., real-time manipulation of graphical images which are created from previously computed datasets.

It is probably realistic to assume that the demand for supercomputers and graphical displays expands to fill any available resource. However,

because of the distributed processing capability of the NPSN and its uniform software environment and protocol suite, coprocessing between workstation and supercomputers can be easily implemented for more effective utilization of available resources.

Graphical coprocessing tasks, such as the ARC-developed Realtime Interactive Particle Traces (RIP) software package which interactively calculates particle paths in a computed flow field, involve initiating two coupled processes simultaneously: a computational and memory intensive process residing on the Cray-2 and a graphical display and manipulation process residing on the workstation. Information is passed back and forth between the two machines via HYPERchannel and Ethernet links for remote workstations. In this way, a processing task can be intelligently divided to make use of the strong points of the various machines in the NAS system. For the user, this translates into a substantial savings in time for the display of a particle trace on a workstation, since the Cray-2/IRIS combination offers a 10:1 performance improvement over an IRIS alone. Because of the uniform UNIX environment across the Cray-2 and workstations, the creation and installation of RIP required one afternoon to coordinate among three people with existing but decoupled codes.

Quality film output of workstation-generated displays is provided by Model 655 Dunn cameras and prototype Seiko printers. Users can make high quality photographs, transparencies, or negatives with the Dunn cameras. The Dunn system works directly off the RGB output of the workstations. For 16-mm movies, the computer-controlled Model 632 Dunn camera functions with the aid of a workstation patch panel.

Video animation capabilities are provided by an ABEKAS digital disk recorder. This system takes an RS170 output from the IRIS, converts it to NTSC signals, and records it on disk. This represents a significant improvement in recording speed, editing, and special effects capabilities over older tape-recording devices. To assist the user in making animation sequences, a program called the Graphics Animation System (GAS) has been developed. GAS lets the user create scripts on the IRIS which can be played back at a later time on the animation system. GAS is a fairly "user friendly" program in that it records a user's actions as graphic images are manipulated on the IRIS screen, automatically generating a playback script. If needed, the user can edit this script to either correct the errors or to add steps to smooth out the animation.

#### FUTURE CHALLENGES AND PLANS

The NAS Program is engaged in planning the implementation of the EOC and its follow-on con-

figurations. EOC, planned to be operational in 1988, adds an advanced supercomputer, the HSP-2, with performance and memory capacity targeted at four times that of the Cray-2.

The addition of this large increment in computer power raises the potential for imbalance among processing rate, data transport bandwidth, and storage capacity across the entire NPSN. At present, the performance and storage capacity of supercomputers is outpacing the ability to move data to and from the computer and to store it. New storage technologies such as optical, magneto-optic, and magnetic vertical recording look promising. Unfortunately, supercomputing needs alone are not a sufficient market force to drive these technologies to the desired level of performance and capacity. As a consequence, new strategies with older technologies are being investigated to meet near-future needs.

Increased high-speed processor data production, higher graphics data consumption, and higher storage capacity will require increased data communication bandwidth in the NPSN. An analysis of the bandwidth requirement for a 20-fold increase in processing power, based on models of the NPSN projected workload and system architecture, shows an effective bandwidth of 100 megabits/sec is necessary (Levin et al., 1987). To meet this requirement, the NAS Program has efforts under way to establish a network test bed for hands-on investigation of new networking technologies.

Long-haul data communications is also an important area for future improvement. The experience gained with NASnet has demonstrated the value of high-bandwidth and low-delay terrestrial circuits. The NAS Program plans to aggressively pursue opportunities to economically increase bandwidth capability when high capacity transcontinental fiber-optic data trunks become available in the next few years. In addition, an advanced remote-communications gateway prototype is planned. The objective of this effort is to meter circuit capacity as throughput demand rises and falls. This will be accomplished by a type-of-service routing gateway that establishes the optimum bandwidth and delay characters to match a particular application service (interactive character stream, on-demand data transfer, bulk file transfer, etc.). The bandwidth will be optimized by adding or deleting subchannel allocations of the backbone circuit and the delay optimized by switching to either low-delay terrestrial or high-delay satellite circuits.

The first-order software challenge for the future is to provide the algorithms, languages, and programming tools to effectively use the potential of parallel architecture computers. Future high-speed processors, without a doubt, will consist of several if not many processors. Parallel supercomputers are here now in the form of the four-processor Cray XMP and Cray-2, but their parallel multitasking capability is greatly



under utilized. When the multitasking of many processors becomes necessary to achieve greater performance, much improved software will be critically needed. The NAS Program is placing greater emphasis on improved techniques for multitasking and will aggressively pursue research in massively parallel architectures.

Future technology will also support more parallelism across computer boundaries as in the case of coprocessing between supercomputers and workstations. Special-purpose processors, designed to perform very well on a piece of the total problem, connected to very fast networks are a fertile area for research and development. The NAS Program is actively engaged in research and advance development in these areas, with the goal of introducing them into the NPSN by the end of the decade.

#### SUMMARY

The NAS program has met its first major milestone--the NPSN IOC. The program has met its goal of providing a national supercomputer facility capable of greatly enhancing the Nation's research and development efforts. Furthermore, the program is fulfilling its pathfinder role by defining and implementing a new paradigm for supercomputing system environments. The IOC is only the begin-

ning and the NAS Program will aggressively continue to develop and implement emerging supercomputer, communications, storage, and software technologies to strengthen computations as a critical element in supporting the Nation's leadership role in aeronautics.

#### REFERENCES

- Jennings, D. M.; Landweber, L. H.; Fuchs, I. H.; Farber, D. J.; and Adrion, W. R.: Computer Networking for Scientists, Science, Vol. 231, Feb. 1986.
- Levin, E.; Eaton, C. K.; and Young, B.: Scaling of Data Communications for Advanced Supercomputer Network. To be pres. to the IFIP/IEEE/ITC Intern. Conf. Data Commun. Systems and Their Performance, Rio de Janeiro, June 22-25, 1987.
- Peterson, V. L.; and Arnold, J. O.: Impact of Supercomputers on Experimentation: A View from a National Laboratory, Proc. Amer. Soc. Eng. Symp., Atlanta, June 17-19, 1985.
- Peterson, V. L.; Ballhaus, W. F., Jr.; and Bailey, F. R.: Numerical Aerodynamic Simulation (NAS), in Large Scale Scientific Computation (Seymour V. Parter, ed.), Academic Press, New York, 1984.